
Routing Metrics and Protocols for Wireless Mesh Networks

Miguel Elias M. Campista, Pedro Miguel Esposito, Igor M. Moraes, Luís Henrique M. K. Costa, and Otto Carlos M. B. Duarte, Universidade Federal do Rio de Janeiro
Diego G. Passos, Célio Vinicius N. de Albuquerque, and Débora Christina M. Saade, Universidade Federal Fluminense
Marcelo G. Rubinstein, Universidade do Estado do Rio de Janeiro

Abstract

WMNs are low-cost access networks built on cooperative routing over a backbone composed of stationary wireless routers. WMNs must deal with the highly unstable wireless medium. Therefore, the design of algorithms that consider link quality to choose the best routes are enabling routing metrics and protocols to evolve. In this work, we analyze the state of the art in WMN metrics and propose a taxonomy for WMN routing protocols. Performance measurements for a WMN, deployed using various routing metrics, are presented and corroborate our analysis.

Wireless networks are becoming increasingly popular as they provide flexibility, mobility support, and are easy to deploy. In addition, the reduced wired infrastructure combined with large-scale commercialization, notably of IEEE 802.11, results in plummeting costs. Thus, more and more Internet service providers (ISPs) offer wireless access that in the long term will result in ubiquitous Internet.

Infrastructure-based wireless networks, such as IEEE 802.11 wireless distribution systems, limit the coverage to users within the transmission range of access points. In this case, access points are connected to a wired network, which incurs high infrastructure costs. Ad hoc networks [1] have no infrastructure costs because they do not require wires. Nevertheless, ad hoc networks cannot supply backhaul access and may become a collection of isolated networks due to user mobility. Choosing the position of access points in wireless distribution systems or predicting the location of a user to avoid isolated areas is challenging.

Wireless mesh networks (WMNs) [2] aim at guaranteeing connectivity. WMNs build a multihop wireless backbone to interconnect isolated local area networks (LANs) and to extend backhaul access to users not within range of typical access points. Backbone routers are usually stationary, and mobile users roam among them. Consequently, they can be permanently power-supplied. As mobility and energy-saving are no longer issues, WMN routing considers link quality metrics, such as capacity or error probability.

Currently, much effort is expended on the IEEE 802.11 medium access control (MAC) layer to fully exploit novel physical (PHY) layer techniques. Nevertheless, in multihop scenarios, performance depends on the routing protocol to properly choose routes, given the current network conditions.

Different metrics and protocols are proposed to improve wireless mesh routing. Additionally, the upcoming IEEE 802.11s defines multihop forwarding at the link layer, making

a WMN appear as a LAN for layer 3 protocols. In this article, we review ongoing research on WMN routing and present performance results obtained with different metrics in our WMN testbed. First, we review state-of-the-art routing metrics. Then, we analyze WMN routing protocols and propose a taxonomy based on their algorithms.

This article is organized as follows. We describe the main WMN routing metrics and protocols. We compare different metrics using our testbed. We conclude this article and identify open research directions.

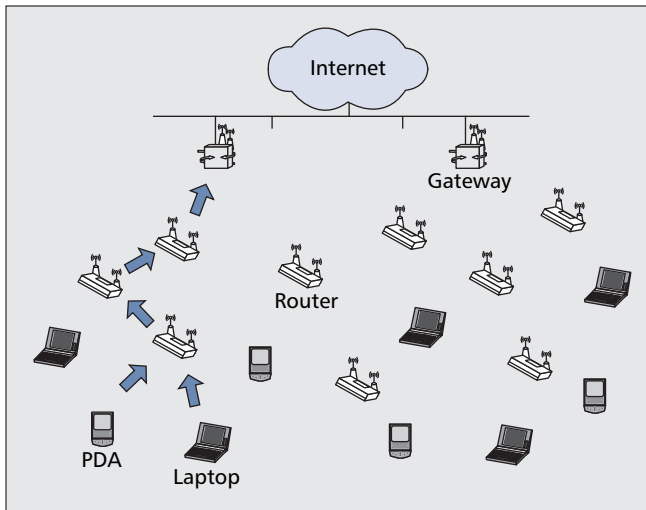
Wireless Mesh Routing

WMN backbone routers use multihop communication similarly to ad hoc networks (Fig. 1). On the other hand, mobile users connect to the backbone via mesh routers that play the role of access points. The backbone routers typically are stationary, which permits routing metrics to model link quality instead of simply using the number of hops. Assuming that the common-case application in WMNs is Internet access, traffic is concentrated on links close to the gateways.

Routing Metrics

Ad hoc networks usually use the hop count as a routing metric. This metric is appropriate for ad hoc networks because new paths must be found rapidly, whereas high-quality routes may not be found in due time. This is important in ad hoc networks because of user mobility. In WMNs, the stationary topology benefits quality-aware routing metrics [3].

The first metric proposed for WMNs is the expected transmission count (ETX) [4]. ETX is the expected number of transmissions a node requires to successfully transmit a packet to a neighbor. To compute ETX, each node periodically broadcasts probes containing the number of received probes from each neighbor. The number of received probes is calculated at the last T time interval in a sliding-window fashion. A



■ Figure 1. A typical wireless mesh network.

node A computes the ETX of the link to a node B by using the delivery ratio of probes sent on the forward (d_f) and reverse (d_r) directions. These delivery ratios are, respectively, the fraction of successfully received probes from A announced by B and the fraction of successfully received probes from B, at the same T interval. The ETX of link AB is $1/(d_f \times d_r)$. The ETX computation considers both forward and reverse directions because of data- and ACK-frame transmission. The chosen route is the one with the lowest sum of ETX along the route to the destination. The number of broadcast probes in an n -node network is $O(n)$. The minimum loss (ML) metric [5] also is based on probing to compute the delivery ratio. Rather than calculating ETX, ML finds the route with the lowest end-to-end loss probability. Thus, ML is not additive as ETX is. Instead, ML multiplies the delivery ratios of the links in the reverse and forward directions to find the best path. The authors of ML argue that the use of multiplication reduces the number of route changes, improving network performance.

The implementation of ETX has revealed two shortcomings: broadcasts usually are performed at the network basic rate, and probes are smaller than typical data packets. Thus, unless the network is operating at low rates, the performance of ETX becomes low because it neither distinguishes links with different bandwidths nor considers data-packet sizes. To cope with these issues, the expected transmission time (ETT) [4] is the time a data packet requires to be transmitted successfully to each neighbor. ETT adjusts ETX to different PHY rates and data-packet sizes.

Currently, there are two main approaches to compute ETT. For Draves *et al.* [4], ETT is the product between ETX and the average time a single data packet requires to be delivered ($ETT = ETX \times t$). To calculate this time t , the authors divide a fixed data-packet size (S) by the estimated bandwidth (B) of each link ($t = S/B$). The authors prefer to periodically estimate the bandwidth than to use rates retrieved from firmware. The packet-pair technique then is used to calculate B per link. This technique consists of transmitting a sequence of two back-to-back packets to estimate bottleneck bandwidth. In the implementation of Draves *et al.*, two packets are unicast in sequence, a small one followed by a large one, to estimate the link bandwidth to each neighbor. Each neighbor measures the inter-arrival period between the two packets and reports it back to the sender. The computed bandwidth is the size of the large packet of the sequence divided by the minimum delay received for that link. In an n -node network where each node has v adjacencies, estimating the bandwidth is $O(n \cdot v)$. Another approach to compute ETT is considered in [6]. The author estimates the

loss probability by considering that IEEE 802.11 uses data and ACK frames. The idea is to periodically compute the loss rate of data and ACK frames to each neighbor. The former is estimated by broadcasting a number of packets of the same size as data frames, one packet for each data rate defined in IEEE 802.11. The latter is estimated by broadcasting small packets of the same size as ACK frames and sent at the basic rate that is used for ACKs. Note that broadcasting packets at higher data rates may require firmware modifications. ETT is the inverse of the product between the best throughput achievable (r_t) and the delivery probability of ACK packets in the reverse direction (p_{ACK}). Computing ETT in an n -node network is $O(n \cdot m)$, where m is the number of possible data rates. Similarly to ETX, the chosen route is the one with the lowest sum of ETT values.

Cross-layer approaches are receiving special attention in WMNs [2]. Among the available techniques, the use of multiple channels is commonplace. Through multiple channels, it is possible to improve network throughput by using, at the same time, the available non-overlapping channels defined by IEEE 802.11. This technique, however, must deal with two issues to become effective, namely, intra-flow and inter-flow interference. The intra-flow interference occurs when different nodes transmitting packets from the same flow interfere with each other. Maximizing the number of channels is not trivial, considering that nodes must maintain connectivity. The inter-flow interference otherwise is the interference suffered among concurrent flows. The weighted cumulative ETT (WCETT) [4] changes ETT to also consider intra-flow interference. This metric is a sum of end-to-end delay and channel diversity. A tunable parameter is used to combine both components or prioritize one of them. Unlike ETX and ETT, WCETT is an end-to-end metric. Thus, its outcome is the final cost of the route. This metric computes end-to-end values because it must consider all channels used along the route to avoid intra-flow interference. Nevertheless, WCETT neither guarantees shortest paths nor avoids inter-flow interference [7]. Link-state-based routing protocols require minimum-cost routes to be loop-free. Moreover, not avoiding inter-flow interference may lead WCETT to choose routes in congested areas. The metric of interference and channel-switching (MIC) addresses these issues [7]. First, each node takes into account the number of interfering nodes in the neighborhood to estimate inter-flow interference. In addition, MIC uses virtual nodes to guarantee the minimum-cost routes computation. MIC also calculates its value based on the ETT metric.

One critical problem of wireless networks is the fast link-quality variation. Metrics based on average values computed on a time-window interval, such as ETX, may not follow the link-quality variations or may produce prohibitive control overhead. Especially in indoor environments, this problem is even more difficult. To cope with this, modified ETX (mETX) and effective number of transmissions (ENT) were proposed [3]. These metrics consider the standard deviation in addition to link-quality average values to project physical-layer variations onto routing metrics.

The mETX metric also is calculated by broadcasting probes. The difference between mETX and ETX is that rather than considering probe losses, mETX works at the bit level. The mETX metric computes the bit error probability using the position of the corrupted bit in the probe and the dependence of these bit errors throughout successive transmissions. This is possible because probes are composed by a previously known sequence of bits. ENT is an alternative approach that measures the number of successive retransmissions per link considering the variance. ENT also broadcasts probes and limits route computation to links that show an acceptable number of retransmissions according to upper-layer requirements. If a

Metric	Quality-aware	Data rate	Packet size	Intra-flow interference	Inter-flow interference	Medium instability
Hop	×	×	×	×	×	×
ETX	√	×	×	×	×	×
ML	√	×	×	×	×	×
ETT	√	√	√	×	×	×
WCETT	√	√	√	√	×	×
MIC	√	√	√	√	√	×
mETX	√	√	√	×	×	√
ENT	√	√	√	×	×	√
iAWARE	√	√	√	√	√	√

■ Table 1. Main routing metrics characteristics.

link shows a number of expected transmissions higher than the maximum tolerated by an upper-layer protocol (e.g., TCP), ENT excludes this link from the routing computation, assigning to it an infinity metric. Both mETX and ENT are aware of the probe size, therefore the inclusion of the data rate is trivial with the two metrics. Another metric that also considers link-quality variation is interference aware (iAWARE) [8]. This metric uses signal to noise ratio (SNR) and signal to interference and noise ratio (SINR) to continuously reproduce neighboring interference variations onto routing metrics. The iAWARE metric estimates the average time the medium is busy because of transmissions from each interfering neighbor. The higher the interference, the higher the iAWARE value. Thus, unlike mETX and ENT, iAWARE considers intra- and inter-flow interference, medium instability, and data-transmission time.

Although there is an increasing number of routing metrics, a consensus has not been achieved. Up to now, most routing protocol implementations prefer metrics with simpler designs such as ETX or ETT. Table 1 summarizes the main characteristics of the routing metrics discussed.

Routing Protocols

Ad hoc routing protocols are usually proactive, reactive, or hybrid. The proactive strategy operates like classic routing on wired networks. Routers keep at least one route to any destination in the network. Reactive protocols, on the other hand, request a route to a destination only when a node has a data packet to send. If a node does not have data packets to send to a particular destination, the node will never request a route to it.

Many WMN routing protocols use similar strategies. Nevertheless, they are adapted to the peculiarities of WMNs, for example, by using a quality-aware routing metric. We propose a taxonomy for WMN routing protocols with four classes: ad hoc-based, controlled-flooding, traffic-aware, and opportunistic. Each class mainly differs on route discovery and maintenance procedures. In WMNs, most routing protocols consider that the network is only composed by wireless backbone nodes. If, eventually, a mobile device operates as a backbone node, it must run the same routing protocol.

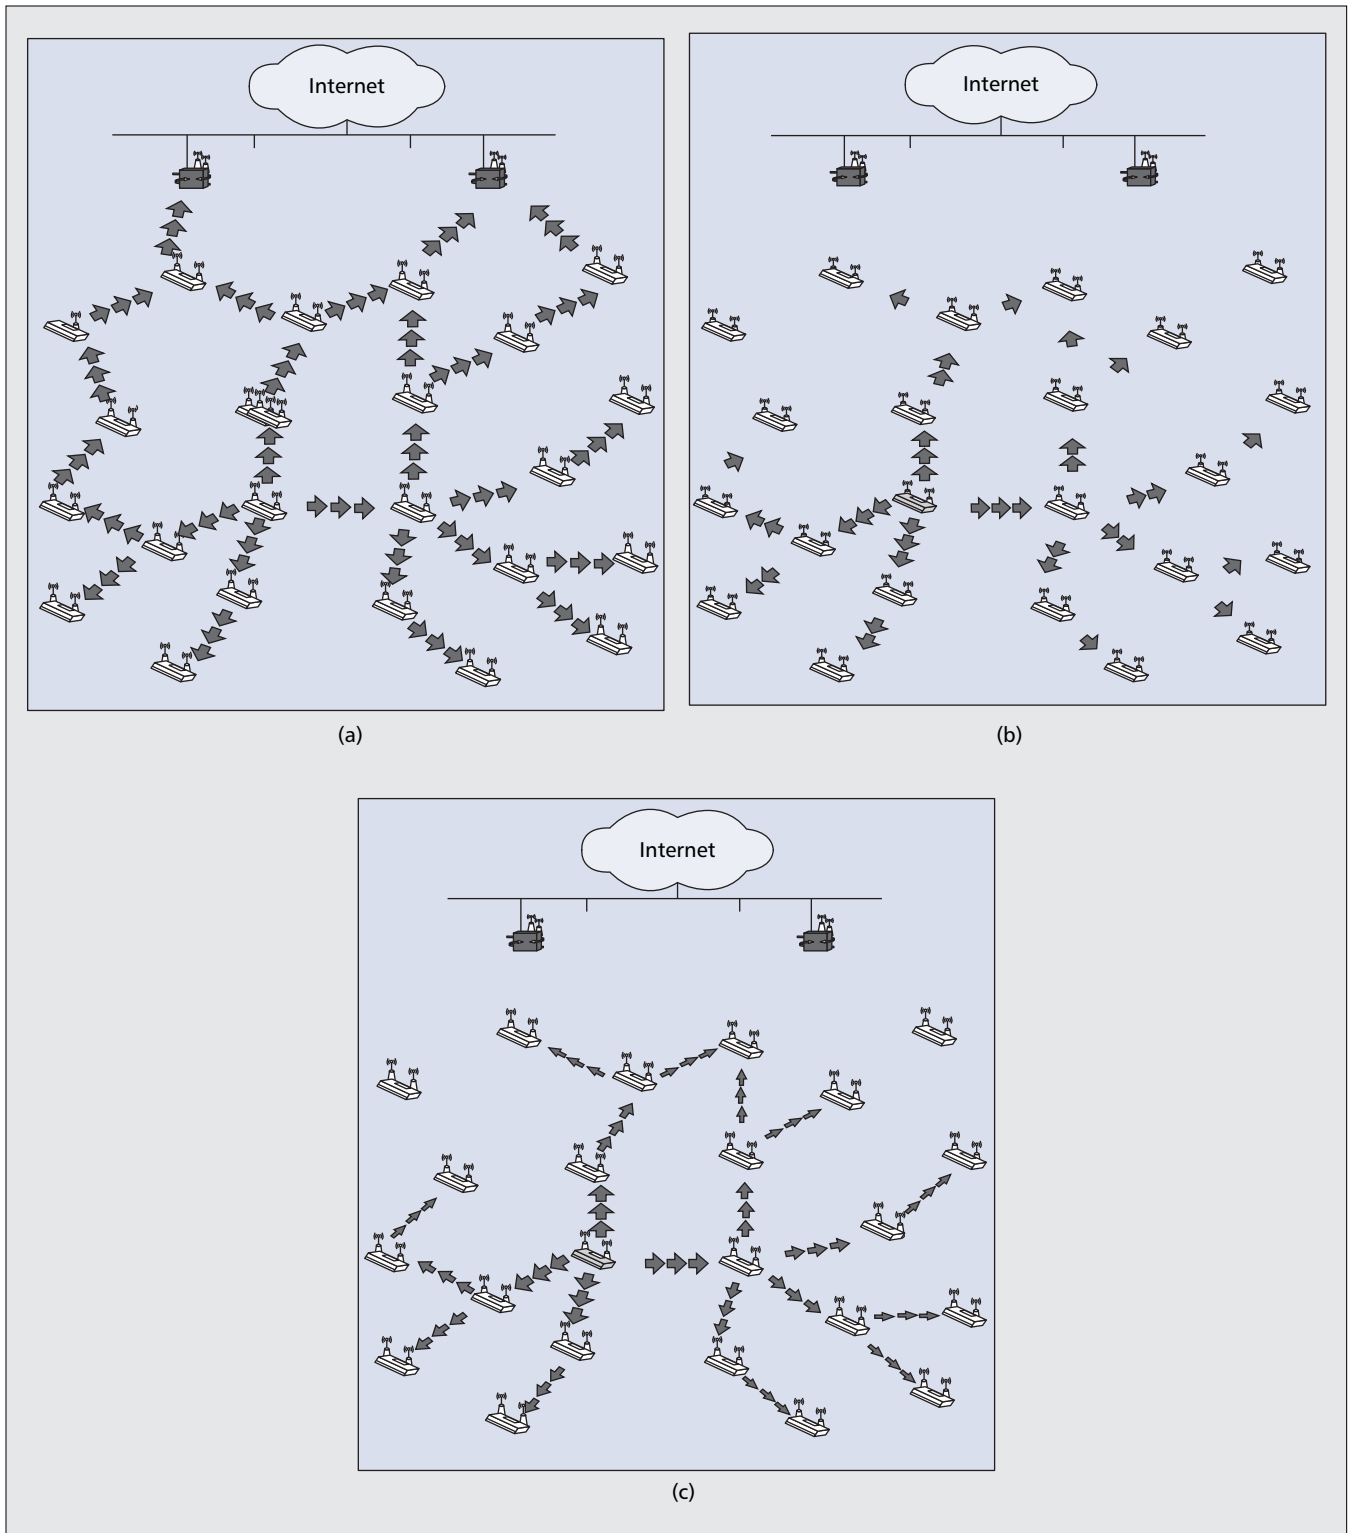
WMN ad hoc-based protocols adapt ad hoc routing protocols to deal with link-quality variations. Routers continuously update their outgoing-link metrics and disseminate them to other routers. The link quality source routing (LQSR) protocol [4] combines link-state proactive routing with the reactive

strategy from ad hoc networks. As a link-state routing protocol, LQSR uses a complete view of the network topology to compute shortest paths. Nevertheless, LQSR uses a route discovery procedure as in reactive protocols to reduce routing overhead, which may become high because of medium instability and user mobility. During route discovery, LQSR obtains up-to-date link state information of the traversed links, reducing the periodicity of regular link-state advertisements. SrcRR [6] is another ad-hoc-based protocol. It uses only a discovery procedure similar to reactive protocols to update the routing information of the traversed links, reducing control overhead. Nevertheless, it computes routes using a reduced view of the network. Both LQSR and SrcRR implement route discovery procedures using source routing and ETX.

Physical-layer techniques usually are used to improve the overall efficiency of routing protocols. The multi radio LQSR (MR-LQSR) [4] adapts LQSR to operate over multiple channels and multiple interfaces, using the WCETT metric. Although WCETT does not guarantee minimum cost paths, MR-LQSR is loop-free because it uses source routing.

Controlled-flooding protocols use algorithms designed to reduce control overhead. Flooding the network with routing updates may produce scalability issues, especially if frequent changes on medium conditions are considered. We identify two baseline approaches that reduce the routing overhead as compared to classical flooding (Fig. 2a). In temporal flooding (Fig. 2b), the periodicity is set according to the distance from the source router. On the other hand, using spatial flooding (Fig. 2c), the distant nodes receive less precise or less detailed information from the source. In practice, most protocols disseminate local-scope routing information, using the temporal approach. The basic assumption is that flooding the network is not efficient because most communication in wireless networks is between nearby nodes. Therefore, there is no need to send control packets to the distant nodes as frequently as to nearby ones. Another way to reduce overhead is to limit the number of nodes responsible for flooding the network, reducing redundancies. A common approach is to use algorithms that find the minimum set of nodes required to forward routing information to all destinations in the network.

The localized on-demand link state (LOLS) [9] attributes a long-term cost and a short-term cost to links. Long-term and short-term costs represent the usual and the current cost of a link, respectively. To reduce control overhead, short-term costs are frequently sent to neighbors, whereas long-term costs are sent using longer periods. LOLS computes routes using ETX or ETT. Another typical example is the Mobile Mesh Routing Protocol (MMRP) developed by the MITRE Corporation. MMRP assigns an age to routing messages as the open shortest path first (OSPF) protocol does. Whenever a node sends a routing message, it subtracts the age of the message by the estimated time required to forward it. Upon age expiration, the message is dropped, preventing its retransmission. MMRP does not specify a routing metric. The optimized link state routing (OLSR) is another example of a controlled-flooding protocol (RFC 3626). OLSR was adapted to use ETX as a link metric in WMNs. It uses the fraction of HELLO messages lost in a given interval of time to calculate ETX. OLSR also could be classified as an ad hoc-based protocol; however, it uses multipoint relays (MPRs), a controlled-flooding technique. OLSR limits the number of nodes in charge of disseminating control packets to reduce redundancies. Each node selects its MPR

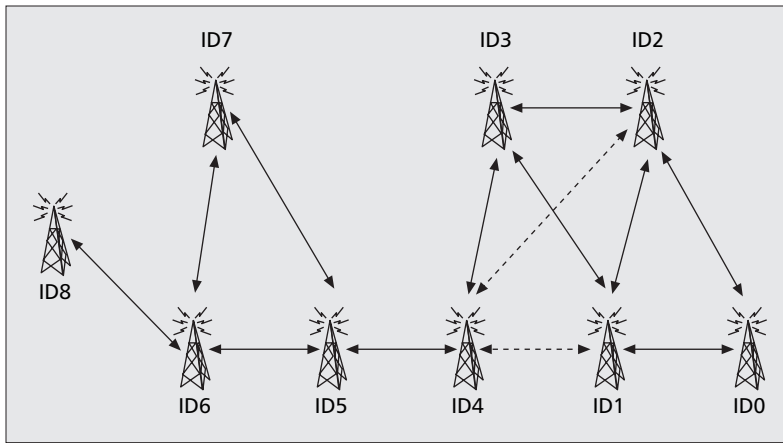


■ Figure 2. Flooding types: a) classical; b) temporal; c) spatial.

set, which is composed of nodes responsible for forwarding routing information from the selector node. Each node constructs an MPR set with the minimum number of one-hop neighbors required to reach all two-hop neighbors.

Traffic-aware, or tree-based protocols, consider the usual traffic matrix of WMNs. Assuming that backhaul access is the common-case application, they consider a tree-like network topology. The ad hoc on-demand distance vector-spanning tree (AODV-ST) [10] adapts the AODV protocol from ad

hoc networks. In AODV-ST, the gateway periodically requests routes to every node in the network to update its routing table. The gateway is the root of the tree. Communications that do not include the gateway use the original AODV. AODV-ST supports ETX and ETT metrics. Raniwala and Chiueh propose a routing algorithm [11] based on the spanning tree used in wired networks. Route maintenance is done with join and leave requests. This protocol uses the hop metric and other metrics for load-balancing.



■ Figure 3. UFF's mesh network.

Opportunistic protocols improve classical routing based on cooperative diversity schemes. Classical routing protocols compute a sequence of hops to the destination before sending a data packet, either using hop-by-hop or source routing. In case of link failures, successive link-layer retransmissions are performed until successful reception at the next-hop neighbor or until the maximum number of link-layer retransmissions is reached. This approach may incur high delay and poor performance because wireless links require time to recover from failures. Cooperative diversity schemes, on the other hand, exploit the broadcast nature of radio-frequency transmission to set multiple paths towards a destination. The receiver requires suitable transceivers to choose one of the relayed signals or to use a combination of them. Opportunistic protocols adapt cooperative diversity to standard IEEE 802.11 transceivers. Therefore, only one node forwards each packet. For example, opportunistic protocols choose, on-the-fly, the next hop that offers the best throughput. These protocols guarantee that the data is always forwarded whenever there is at least one next hop. In addition, the chosen route likely uses the best quality links, considering short-term variations.

The ExOR protocol combines routing with MAC layer functionality [12]. Routers send broadcast packets in batches, with no previous route computation. Packets are transmitted in batches to reduce protocol overhead. In addition, broadcasting data packets improves reliability because only one intermediate router is required to overhear a transmission. Nevertheless, it does not guarantee that packets are received because they are not acknowledged. Thus, an additional mechanism is required to indicate correct data reception. Among the intermediate routers that have heard the transmission, only one retransmits at a time. The source router defines a forwarding list and adds it to the header of the data packets. This list contains the addresses of neighbors, ordered by forwarding priority. Routers are classified in the forwarding list according to their proximity to the destination, computed by a metric similar to ETX. The metric used by ExOR considers only the loss rate in the forward direction because there are no acknowledgments. Upon reception of a data packet, the intermediate router checks the forwarding list. If its address is listed, it waits for the reception of the whole batch of packets. It is possible, however, that a router does not receive the entire batch. To cope with this problem, the highest-priority router that has received packets forwards them and indicates to the lower-priority routers the

packets that were transmitted. Consequently, the lower-priority routers transmit the remaining packets, avoiding duplicates. The transmissions are performed until the destination indicates the reception. The Resilient Opportunistic Mesh Routing protocol (ROMER) [13] combines long-term shortest-path or minimum-latency routes with on-the-fly opportunistic forwarding to provide resilient routes and to deal with short-term variations on medium quality. ROMER computes long-term routes and opportunistically expands or shrinks them at run time to fully exploit short-term higher-quality links. Long-term routes are computed using the minimum number of hops or the minimum average delay. Unlike ExOR, ROMER transmits on a packet basis to enable faster reaction to medium variations.

The highest-throughput route is chosen according to the maximum PHY rate as indicated by the MAC layer.

Table 2 presents the main routing protocols according to our taxonomy and lists the main routing metrics used by each protocol.

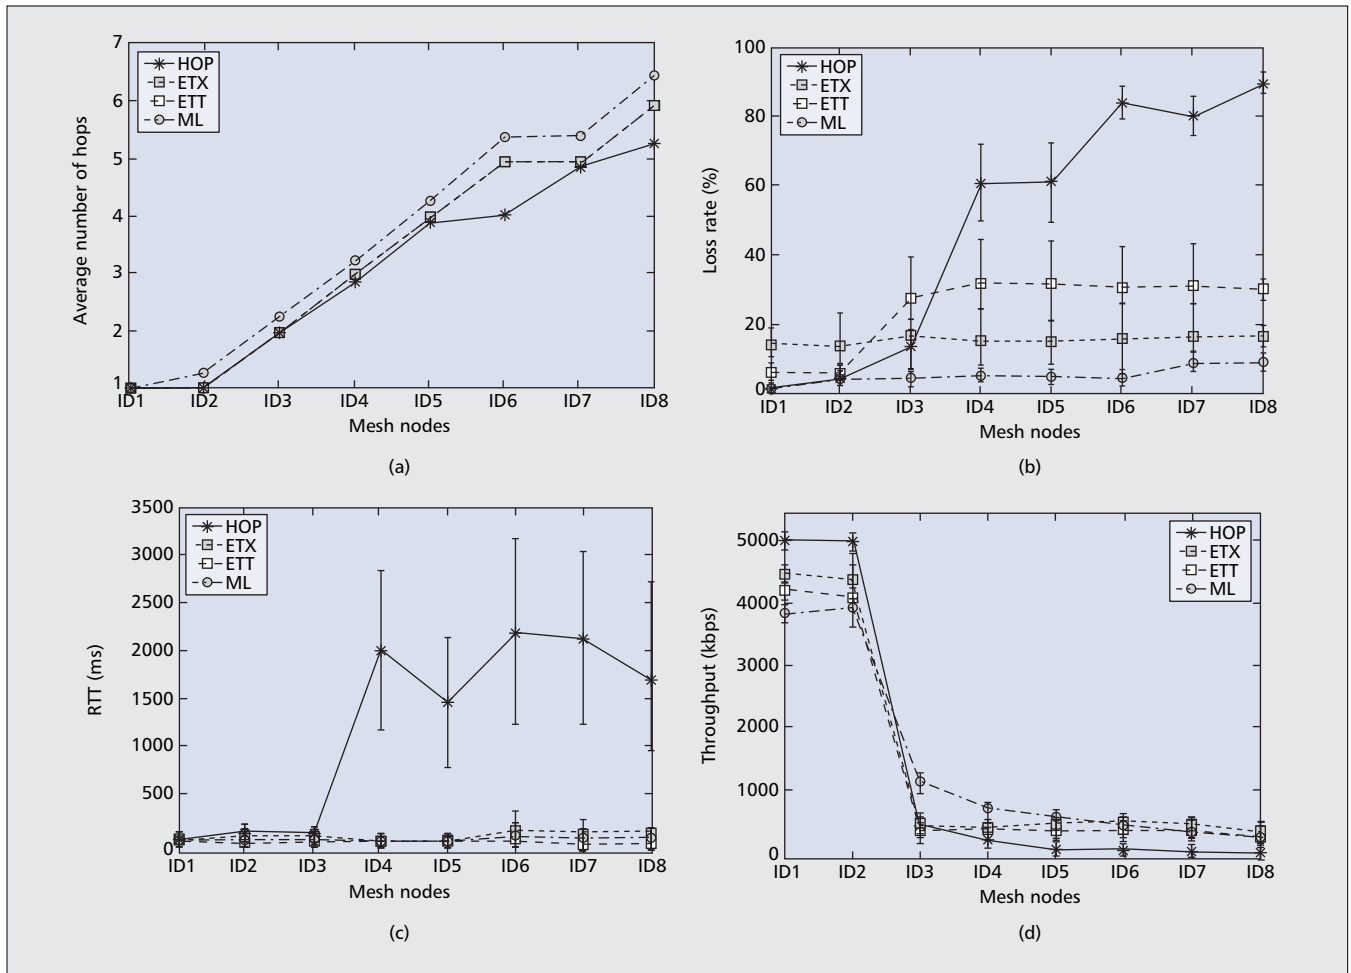
Mesh Network Performance Analysis

This section evaluates the performance of different WMN routing metrics. Hop count, ETX, ETT, and ML metrics are implemented and assessed using the OLSR routing protocol. The link-state-based routing protocol OLSR is being defined by the upcoming IEEE 802.11s standard as the basis for future routing protocol implementations defined at the link layer.

Our performance measurements were collected in the ReMesh mesh network deployed at the Fluminense Federal University (UFF) campus in the city of Niterói, Brazil. Measurements were performed in the indoor testbed using programmable wireless routers based on the OpenWRT open-source operating system. These routers are Linksys WRT54G/GS/GL 802.11g using their native 2 dB omni-directional antennas. The mesh network deployed at UFF consists of nine mesh nodes labeled from ID0 to ID8 deployed at the third and fourth floors of the engineering building of the university (Fig. 3). Node IDs are num-

Class	Protocols	Metrics
Ad hoc based	LQSR	ETX
	SrcRR	ETX
	MR-LQSR	WCETT
Controlled flooding	LOLS	ETX or ETT
	MMRP	Not specified
	OLSR	Hop, ETX, ML, or ETT
Traffic-aware	AODV-ST	ETX or ETT
	Raniwala and Chiueh's	Hop or load-balancing metrics
Opportunistic	ExOR	Unidirectional ETX
	ROMER	Hop or delay

■ Table 2. WMN protocols and their respective routing metrics.



■ Figure 4. Performance results for hop, ETX, ETT, and ML metrics: a) average route length; b) packet loss rate; c) round trip-time; d) throughput.

bered according to their physical distance to node ID0. Wireless links connecting nodes were collected by monitoring the topology built by OLSR within each router, using a plug-in for the OLSR daemon. Dashed lines indicate low quality links with loss rates higher than 50 percent, and continuous lines indicate better quality links. The OLSR daemon natively implements hop and ETX metrics; we implemented ETT and ML. In the ML case, we changed the OLSR implementation to use multiplicative metrics instead of additive ones. In the ETT case, we developed a plug-in for the OLSR daemon to calculate ETT according to the packet-pair technique [4].

Number of Hops

Figure 4a shows the average number of hops traversed to reach each node from node ID0 for each metric. It can be observed that on average, using the hop metric, each node is reached with the lowest number of hops, whereas the ML metric chooses paths with the highest number of hops. ETX and ETT tend to select routes with the same number of hops, but not necessarily the same route. Results are consistent with the physical distance between the nodes and with the quality of the links between them (Fig. 3).

Packet Loss Rate

To evaluate the packet loss rate (PLR) experienced when using each routing metric, an experiment was performed over a 24-hour period, transmitting in each run 600 ping packets between node ID0 and every other node of the network. Each run was repeated 36 times for each of the four metrics in a round-robin fashion.

Figure 4b shows the average packet loss rate experienced at each node ID for each metric. All measurements are presented with a confidence interval of 90 percent. As the distance to node ID0 grows, the use of the hop metric results in increasingly high packet-loss rates. This behavior is expected because the hop metric does not consider the quality of the links and tends to forward packets through long noisy wireless links. ETX and ETT metrics converged to PLR in the order of 19 percent and 30 percent, respectively, regardless of the distance to node ID0. The ML metric performed best among the four metrics, because it is designed to select routes with low loss links. The ML metric resulted in PLR in the range of 5 percent for up to node ID6 and around 10 percent for nodes ID7 and ID8.

Network Delay

During the PLR experiment, the average round-trip-time (RTT) for packets traveling from node ID0 to each other node and back was also collected (Fig. 4c). All measurements are presented with a confidence interval of 90 percent. As the distance to node ID0 grows, the use of the hop metric results in high RTTs on the order of two seconds. This behavior occurs because, although the route taken when using the hop metric has a smaller number of hops, the noisy links used by this metric result in a high number of layer-2 retransmissions and therefore, in longer delays to forward layer-3 packets. All other metrics achieved RTTs lower than 150 ms for ETX, 75 ms for ML, and 35 ms for ETT. The ETT metric is the only one to estimate the transmission time, and this feature produced the best performance in terms of RTT.

Throughput

To evaluate the throughput experienced when using each routing metric, an experiment was performed over a 24-hour period, performing, in each run, a total of 600 IPERF-TCP measurements between node ID0 and every other node of the network. Each run was repeated 36 times for each of the four metrics in a round-robin fashion.

Figure 4d shows the average throughput in kb/s experienced at each node ID for each metric. All measurements are presented with a confidence interval of 90 percent. This experiment is interesting because typically ETX, ETT, and ML choose paths with a higher number of hops when compared to the hop metric. Each additional hop in multihop transmissions over the shared medium increase contention and collision probability and can have a negative impact on throughput. For short distances, all metrics achieved high throughput with hop leading to throughputs in the order of 5 Mb/s. As the distance increased, the hop metric throughputs dropped significantly to close to zero, whereas all other metrics exhibited similar performance resulting on throughputs in the order of 500 kb/s.

Conclusion

In this article, we reviewed the main WMN routing metrics and proposed a taxonomy for the main WMN routing protocols. We showed that the evolution of quality-aware metrics requires an incremental complexity in metric computation.

Routing protocols were classified in four categories: ad hoc-based, traffic-aware, controlled-flooding, and opportunistic. All protocols aim to better utilize wireless medium resources but use different approaches, such as mixing reactive and proactive strategies, considering tree-based approximations of the network topology, reducing control overhead, or increasing medium access reliability. All of these control dissemination techniques can be combined with the proposed quality-aware link metrics.

We also showed performance measurements collected in the ReMesh mesh network deployed at the UFF campus in Niterói, Brazil. We tested the performance of four metrics, assessed using the OLSR protocol, namely hop, ETX, ML, and ETT. Our results confirmed that the hop metric performs poorly because it is not aware of link-quality variations. On the other hand, ML, ETX, and ETT, showed better results, considering the different performance measures in accordance with the design of each metric.

The design of WMNs presents a number of open issues, ranging from routing metrics to security. One direction is cross-layer design to improve routing efficiency. This is accomplished by better reflecting PHY-layer variations onto routing metrics or by better using the available radio spectrum to directly improve the network throughput.

Acknowledgments

We would like to thank CNPq, CAPES, FAPERJ, FUNTEL, and FUJB. Also, we would like to thank Felipe Schiller for his contribution in the ETT implementation.

References

- [1] M. E. M. Campista *et al.*, "The Ad Hoc Return Channel: A Low-Cost Solution for Brazilian Interactive Digital TV," *IEEE Commun. Mag.*, vol. 45, no. 1, Jan. 2007, pp. 136–43.
- [2] I. F. Akyildiz and X. Wang, "A Survey on Wireless Mesh Networks," *IEEE Commun. Mag.*, vol. 43, no. 9, Sept. 2005, pp. S23–S30.
- [3] C. E. Koksal and H. Balakrishnan, "Quality-Aware Routing Metrics For Time-Varying Wireless Mesh Networks," *IEEE JSAC*, vol. 24, no. 11, Nov. 2006, pp. 1984–94.
- [4] R. Draves, J. Padhye, and B. Zill, "Routing in Multi-Radio, Multi-Hop Wireless Mesh Networks," *ACM MobiCom*, Sept. 2004, pp. 114–28.

- [5] D. Passos *et al.*, "Mesh Network Performance Measurements," *Int'l. Info. and Telecommun. Technologies Symp.*, Dec. 2006.
- [6] D. S. J. de Couto, "High-Throughput Routing for Multi-Hop Wireless Networks," Ph.D. diss., MIT, 2004.
- [7] Y. Yang, J. Wang, and R. Kravets, "Designing Routing Metrics for Mesh Networks," *IEEE Wksp. Wireless Mesh Networks*, Sept. 2005.
- [8] A. P. Subramanian, M. M. Buddhikot, and S. C. Miller, "Interference Aware Routing in Multi-Radio Wireless Mesh Networks," *IEEE Wksp. Wireless Mesh Networks*, Sept. 2006, pp. 55–63.
- [9] S. Nelakuditi *et al.*, "Blacklist-Aided Forwarding in Static Multihop Wireless Networks," *IEEE SECON '05*, Sept. 2005, pp. 252–62.
- [10] K. N. Ramachandran *et al.*, "On the Design and Implementation of Infrastructure Mesh Networks," *IEEE Wksp. Wireless Mesh Networks*, Sept. 2005.
- [11] A. Raniwala and T.-C. Chiueh, "Architecture and Algorithms for an IEEE 802.11-Based Multi-Channel Wireless Mesh Network," *IEEE INFOCOM*, Mar. 2005, pp. 2223–34.
- [12] S. Biswas and R. Morris, "ExOR: Opportunistic Multi-Hop Routing for Wireless Networks," *ACM SIGCOMM*, Aug. 2005, pp. 133–44.
- [13] Y. Yuan *et al.*, "ROMER: Resilient Opportunistic Mesh Routing for Wireless Mesh Networks," *IEEE Wksp. Wireless Mesh Networks*, Sept. 2005.

Biographies

MIGUEL ELIAS M. CAMPISTA [S'05] (miguel@gta.uff.br) received a telecommunications engineering degree from Fluminense Federal University (UFF), Rio de Janeiro, Brazil, in 2003 and an M.Sc. degree in electrical engineering from the Federal University of Rio de Janeiro (UFRJ), Brazil, in 2005. Currently, he is a D.Sc. student with GTA/COPPE/UFRJ. His major research interests are in multihop wireless networks, quality of service, wireless routing, and home networking.

DIEGO G. PASSOS (dpassos@ic.uff.br) received a B.Sc. degree in computer science from UFF in 2007. Currently, he is an M.Sc. student at UFF. His major research interests are multihop wireless networks and wireless routing.

PEDRO MIGUEL ESPOSITO (pedro@gta.uff.br) is a fourth-year undergraduate student in computer engineering at UFRJ. He has been with GTA/COPPE/UFRJ since 2004. His major research interests are in wireless networks, wireless mesh networks, and digital television.

IGOR M. MORAES (igor@gta.uff.br) received an electronic engineering degree and an M.Sc. degree in electrical engineering from UFRJ in 2004 and 2006, respectively. Since March 2006 he has been a D.Sc. student with GTA/COPPE/UFRJ. His major research interests are in optical networks, security, and wireless networks.

CÉLIO VINÍCIUS N. DE ALBUQUERQUE [S'94, M'00] (celio@ic.uff.br) received B.S. and M.S. degrees in electrical and electronics engineering from UFRJ in 1993 and 1995, and M.S. and Ph.D. degrees in information and computer science from the University of California at Irvine in 1997 and 2000, respectively. From 2000 to 2003 he served as the networking architect for Magis Networks, designing high-speed wireless medium access control. Since 2004 he has been an associate professor in the Computer Science Department of UFF. His research interests include Internet architectures and protocols, wireless networks, multicast and multimedia services, and traffic control mechanisms.

DÉBORA CHRISTINA M. SAADE (debora@midia.com.uff.br) received a computer engineering degree, and M.Sc. and D.Sc. degrees in computer science from PUC-Rio, Brazil, in 1992, 1996, and 2003, respectively. Since 2002 she has been an associate professor with the Telecommunications Engineering Department of UFF. Her major research interests are mesh networks, ad hoc routing protocols, QoS, multicast, multimedia communications, multimedia authoring languages, and telemedicine applications.

MARCELO G. RUBINSTEIN (rubi@uerj.br) received a B.Sc. degree in electronics engineering and M.Sc. and D.Sc. degrees in electrical engineering from UFRJ in 1994, 1996, and 2001, respectively. From January to September 2000 he was at the PRISM Laboratory, University of Versailles, France. He is now an associate professor with Universidade do Estado do Rio de Janeiro. His major interests are in wireless networks, home networking, medium access control, and quality of service.

LÚIS HENRIQUE M. K. COSTA [M'99] (luish@gta.uff.br) received his electronic engineer degree and M.Sc. degree in electrical engineering from UFRJ in 1997 and 1998, respectively, and a D.Sc. degree from the University Pierre et Marie Curie (Paris 6), France, in 2001. Since August 2004 he has been an associate professor with COPPE/UFRJ. His major research interests are in the areas of routing, wireless networks, and group communications. He has been a member of the ACM since 2001.

OTTO CARLOS M. B. DUARTE (otto@gta.uff.br) received electronic engineer and M.Sc. degrees from UFRJ in 1976 and 1981, respectively, and a Dr.Ing. degree from ENST/Paris, France, in 1985. Since 1978 he has been a professor with UFRJ. From January 1992 to June 1993 he was with MASI Laboratory, Université Paris 6. In 1995 he spent three months with the International Computer Science Institute (ICSI), University of California, Berkeley. In 1999, 2001, and 2006 he was an invited professor at Université Paris 6. His major research interests are in multicast, QoS guarantees, security, and mobile communications.